




PATENT ^{#3}

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as first class mail in an envelope addressed to Commissioner of Patents and Trademarks, Washington, D.C. 20231 on June 11, 2001.


Signature

Applicant : Mike Wilson, et al.
Application No. : 09/718,150
Filed : November 20, 2000
Title : FLOW INTEGRITY FOR PATH
TRANSITIONING DATA SWITCH
Grp./Div. :
Examiner : Not assigned
Docket No. : 134001/40600

DECLARATION UNDER 37 CFR §1.47
PROOF OF INVENTOR'S REFUSAL TO SIGN

Assistant Commissioner for Patents
Washington, D.C. 20231

26801 West Agoura Road
Calabasas, California 91301
June 11, 2001

Commissioner:

1. I, David A. Cordeiro, am in-house Intellectual Property Counsel for Alcatel Internetworking, Inc. (AII), formerly known as Xylan Corporation ("Xylan").
2. Mr. Chris Hoogenboom is an inventor of the present application, Ser. No. 09/718,150, entitled "Flow Integrity for Path Transitioning Data Switch." Mr. Hoogenboom et al. conceived of the invention disclosed in the present application during his employment at Xylan.

3. On or around January 14, 2000, Mr. Hoogenboom terminated his employment with AII. By virtue of his Employment Agreement with Xylan, now AII, Mr. Hoogenboom agreed in writing to assign to Xylan (AII) all of his rights in and to any inventions he developed while employed by Xylan (AII). Although Mr. Hoogenboom's standard Xylan Corporation Non-Disclosure Agreement is somehow missing from his personnel file, Mr. Hoogenboom's signature on his acceptance letter and continued employment with Xylan, now AII, from September 19, 1994 to January 14, 2000 evidences his reading and agreeing to the standard Xylan Corporation Non-Disclosure Agreement. A copy of Mr. Hoogenboom's acceptance letter and a copy of the standard Xylan Corporation Non-Disclosure Agreement are attached as Exhibits A and B, respectively.

4. On February 18, 2000, the present invention was filed as a Provisional Application (Ser. No. 60/185,153) naming Christopher Hoogenboom et al. as joint inventors.

5. On November 20, 2000, the present invention was filed as a Non-provisional Application claiming the benefit of the aforementioned provisional application and naming Christopher Hoogenboom et al. as joint inventors.

6. On February 16, 2001, Scot Reader (previous Lead Intellectual Property Counsel) and I sent via Federal Express a package containing three documents to Mr. Hoogenboom's litigation counsel, Mr. Eugene L. Hahm of Coudert Brothers, located at 530 Lytton Avenue, Suite 300, Palo Alto, CA 94301. Mr. Hahm is representing Mr. Hoogenboom in a case between AII, Mr. Hoogenboom, Mr. John Wallner (co-inventor of the subject invention), and Internet Machines Corp. The package contained (1) a cover letter, (2) the present application entitled "Flow Integrity for Path Transitioning Data Switch," (3) the application's Assignment and (4) the application's Declaration and Power of Attorney. In regard to the February 16, 2001 correspondence, we received no response

from either Mr. Hoogenboom or Mr. Hahm. Copies of the four documents sent to Mr. Hahm are attached as Exhibit C, and a copy of the FEDEX receipt for the package enclosing the documents is attached as Exhibit D.

7. On or around April 16, 2001, I attempted to contact Mr. Hahm via email regarding the subject application. In regard to the April 16, 2001 correspondence, I received no response from either Mr. Hoogenboom or Mr. Hahm. A copy of the email sent to Mr. Hahm is attached as Exhibit E.

8. On or around April 19, 2001, I sent another correspondence to Mr. Hahm via Federal Express. The letter stated that I had attempted unsuccessfully to contact Mr. Hahm several times regarding his client Mr. Hoogenboom, that I believed Mr. Hoogenboom to be an inventor of the present application, and that as an inventor and employee of AII at the time he conceived of the invention, Mr. Hoogenboom was obligated to sign the Declaration and Assignment and Power of Attorney forms for the present application. Furthermore, I again requested in the letter that Mr. Hahm contact his client, Mr. Hoogenboom, to have him fulfill his duty and sign the Assignment and Declaration and Power of Attorney for the present application. A copy of the letter sent via Federal Express to Mr. Hahm is attached as Exhibit F, and a copy of the FEDEX receipt enclosing the letter is attached as Exhibit G.

9. On April 23, 2001, I received a phone call from Mr. Eugene Hahm's secretary, Ms. Esther Malcher, who said that she received my April 19, 2001 correspondence, that Mr. Hahm was out of town and that a response would be forthcoming. As of today, I have not received any response from either Mr. Hoogenboom or Mr. Hahm in regard to my April 19, 2001 correspondence or any other correspondence.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Respectfully submitted,

DATED: 6/11/01

By: 

David A. Cordeiro

Reg. No. P-48,134

26801 W. Agoura Road

Calabasas, CA 91301

(818) 878-5080



EXHIBIT A



XYLAN

26679 W. Agoura Rd.
Calabasas, CA 91302
☎ (818) 880-3500
☎ (818) 880-3505

July 12, 1994

Mr. Chris Hoogenboom
618 Westbourne Street
La Jolla, CA 92037

Dear Chris:

We are pleased to confirm our offer of employment with Xylan Corporation as Principle Hardware Engineer. Your compensation and benefits are outlined on the attached pages of this letter.

The offer contained in this letter constitutes the entire compensation agreement between you and the Company. Any verbal agreement, assurances or understanding will not alter this offer.

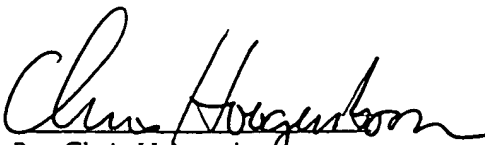
Please sign the copy of this letter indicating your acceptance of this offer and also sign the Compensation Plan indicating acceptance of the Plan. By signing this agreement you understand and agree that either you or Xylan Corporation can terminate the employment relationship at any time with or without cause. Further, no person acting on behalf of Xylan Corporation is authorized to alter this policy, except in a specific written contract. This offer is made contingent upon your reading and signing the standard Xylan Corporation Non-Disclosure Agreement.

We welcome you to the Xylan Corporation team and are looking forward to a mutually successful and beneficial relationship.

Very truly yours,

John Bailey
Director, Engineering

Acceptance:

 9/19/94
By: Chris Hoogenboom Start Date



XYLAN

Chris Hoogenboom

Page 2

Title: Principle Hardware Engineer

Base Location: Chatsworth, California

Position Reports to: John Bailey, Director, Operations

Salary: \$72,000. per year

Relocation: Up to a maximum of \$2,000.00 for relocation expenses when receipts are presented for such expenses.

Review: Annually

Medical/Dental/Life Insurance: Blue Cross, Employee Contribution \$7.00 bi-weekly, dependent coverage \$15.00 bi-weekly.

Vacation: Ten days per year

Stock Options:

Total Option: 20,000

Option Price: \$0.12

Vesting: 5 Years

Please Note: All stock options are subject to approval by the XYLAN Board of Directors, at their next meeting.

Reviewed and Accepted by:


By: Chris Hoogenboom

Date: 7/20/94



EXHIBIT B

XYLAN CORPORATION

EMPLOYMENT, CONFIDENTIAL INFORMATION AND INVENTION ASSIGNMENT AGREEMENT

As a condition of my employment with Xylan Corporation, its subsidiaries, affiliates, successors or assigns (together the "Company"), and in consideration of my employment with the Company and my receipt of the compensation now and hereafter paid to me by Company, I agree to the following:

1. At-Will Employment. I understand and acknowledge that my employment with the Company is for an unspecified duration and constitutes "at-will" employment and any representation to the contrary is unauthorized and not valid unless obtained in writing and signed by an officer of the Company. I acknowledge that this employment relationship may be terminated at any time, with or without good cause or for any or no cause, at the option either of the Company or myself, with or without notice.

2. Confidential Information.

(a) Company Information. I agree at all times during the term of my employment and thereafter, to hold in strictest confidence, and not to use, except for the benefit of the Company, or to disclose to any person, firm or corporation without written authorization of the Board of Directors of the Company, any Confidential Information of the Company. I understand that "Confidential Information" means any Company proprietary information, technical data, trade secrets or know-how, including, but not limited to, research, product plans, product, services, customer lists and customers (including, but not limited to, customers of the Company on whom I called or with whom I became acquainted during the term of my employment), markets, software, developments, inventions, processes, formulas, technology, designs, drawings, engineering, hardware configuration information, marketing, finances or other business information disclosed to me by the Company either directly or indirectly in writing, orally or by drawings or observation of parts or equipment. I further understand that Confidential Information does not include any of the foregoing items which has become publicly known and made generally available through no wrongful act of mine or of others who were under confidentiality obligations as to the item or items involved.

(b) Former Employer Information. I agree that I will not, during my employment with the Company, improperly use or disclose any proprietary information or trade secrets of any former or concurrent employer or other person or entity and that I will not bring onto the premises of the Company any unpublished document or proprietary information belonging to any such employer, person or entity unless consented to in writing by such employer, person, or entity.

(c) Third Party Information. I recognize that the Company has received and in the future will receive from third parties their confidential or proprietary information subject to a duty on the Company's part to maintain the confidentiality of such information and to use it only for certain limited purposes. I agree to hold all such confidential or proprietary information in the strictest confidence and not to disclose it to any person, firm or corporation or to use it except as necessary in carrying out my work for the Company consistent with the Company's agreement with such third party.

3. Inventions.

(a) Inventions Retained and Licensed. I have attached hereto, as Exhibit A, a list describing all inventions, original works or authorship, developments, improvements, and trade secrets which were made by me prior to my employment with the Company (collectively referred to as "Prior Inventions"), which belong to me, which relate to the Company's proposed business, products or research and development, and which are not assigned to the Company hereunder; or, if no such list is attached, I represent that there are no such Prior Inventions. If in the course of my employment with the Company, I incorporate into a Company product, proceeds or machine a Prior Invention owned by me or in which I have an interest, the Company is hereby granted and shall have a nonexclusive, royalty-free, irrevocable, perpetual, worldwide license to make, have made, modify, use and sell such Prior Invention as part of or in connection with such product, process or machine.

(b) Assignment of Inventions. I agree that I will promptly make full written disclosure to the Company, will hold in trust for the sole right and benefit of the Company, will assign and hereby assign to the Company, or its designee, all my right, title, and interest in and to any and all inventions, original works of authorship, developments, concepts, improvements, designs, discoveries, ideas, trademarks or trade secrets, whether or not patentable or registrable under copyright or similar laws, which I may solely or jointly conceive or develop or reduce to practice during the period of time I am in the employ of the Company (collectively referred to as "Inventions"), except as provided in Section 3(f) below. I further acknowledge that all original works of authorship which are made by me (solely or jointly with others) within the scope of and during the period of my employment with the Company and which are protectible by copy right are "works made for hire," as that term is defined in the United States Copyright Act.

(c) Inventions Assigned to the United States. I agree to assign to the United States government all my right, title, and interest in and to any and all Inventions whenever such full title is required to be in the United States by a contract between the Company and the United states or any of its agencies.

(d) Maintenance of Records. I agree to keep and maintain adequate and current written records of all Inventions made by me (solely or jointly with others) during the term of my employment with the Company. The records will be in the form of notes, sketches, drawings, and any other format that may be specified by the Company. The records will be available to and remain the sole property of the Company at all times.

(e) Patent and Copyright Registrations. If the Company is unable because of my mental or physical incapacity or for any other reason to secure my signature on any such document, then I hereby irrevocably designate and appoint the Company and its duly authorized officers and agents as my agent and attorney in fact to act for and in my behalf and stead to execute and file any such document and to do all other lawfully permitted acts to further the prosecution and issuance of letters patent or copyright registrations thereon with the same legal force and effect as if executed by me.

(f) Exception to Assignments. I understand that the provisions of this Agreement requiring assignment of Inventions to the Company do not apply to any invention which qualifies fully under the provisions of California Labor code Section 2870 (attached hereto as Exhibit B). I will advise the Company promptly in writing of any inventions that I believe meet the criteria in California Labor Code Section 2870 and no otherwise disclosed on Exhibit A.

4. Conflicting Employment. I agree that, during the term of my employment with the Company, I will not engage in any other employment, occupation, consulting or other business activity directly related to the business in which the Company is now involved or becomes involved during the term of my employment, nor will I engage in any other activities that conflict with my obligations to the Company.

5. Returning Company Documents. I agree that, at the time of leaving the employ of the Company, I will deliver to the Company (and will not keep in my possession, recreate or deliver to anyone else) any and all devices, records, data, notes, reports, proposals, lists, correspondence, specifications, drawings, blueprints, sketches, materials, equipment, other documents or property, or reproductions of any aforementioned items developed by me pursuant to my employment with the Company or otherwise belonging to the Company, its successors or assigns. In the event of the termination of my employment, I agree to sign and deliver the "Termination Certification" attached hereto as Exhibit C.

6. Notification to New Employer. In the event that I leave the employ of the Company, I hereby grant consent to notification by the Company to my new employer about my rights and obligations under this Agreement.

7. Solicitation of Employees. I agree that for a period of twelve (12) months immediately following the termination of my relationship with the Company for any reason, whether with or without cause, I shall not either directly or indirectly solicit, induce, recruit or encourage any of the Company's employees to leave their employment, or take away such employees, or attempt to solicit, induce, recruit, encourage or take away employees of the Company, either for myself or for any other person or entity.

8. Conflict of Interest Guidelines. I agree to diligently adhere to the Conflict of Interest Guidelines attached as Exhibit D hereto.

9. Representations. I agree to execute any proper oath or verify any proper document required to carry out the terms of this Agreement. I represent that my performance of all the terms of this Agreement will not breach any agreement to keep in confidence proprietary information acquired by me in confidence or in trust prior to my employment by the Company. I have not entered into, and I agree I will not enter into, any oral or written agreement in conflict herewith.

10. Arbitration and Equitable Relief.

(a) Arbitration. Except as provided in Section 10(b) below, I agree that any dispute or controversy arising out of or relating to any interpretation, construction, performance or breach of this Agreement shall be settled by arbitration to be held in Santa Clara County, California, in accordance with the rules then in effect of the American Arbitration Association. The arbitrator may grant injunctions or other relief in such dispute or controversy. The decision of the arbitrator shall be final, conclusive and binding on the parties to the arbitration. Judgment may be entered on the arbitrator's decision in any court having jurisdiction. The Company and I shall each pay one-half of the costs and expenses of such arbitration, and each of us shall separately pay our counsel fees and expenses.

(b) Equitable Remedies. I agree that it would be impossible or inadequate to measure and calculate the Company's damages from any breach of the covenants set forth in Sections 2, 3, and 5 herein. Accordingly, I agree that if I breach any of such Sections, the Company will have available, in addition to any other right or remedy available, the right to obtain an injunction from a court of competent jurisdiction restraining such breach or threatened breach and to specific performance of any such provision of this Agreement. I further agree that no bond or other security shall be required in obtaining such equitable relief and I hereby consent to the issuance of such injunction and to the ordering of specific performance.

11. General Provisions.

(a) Governing Law; Consent to Personal Jurisdiction. This Agreement is made and shall be construed according to the laws of the State of California, without reference to conflict of laws principles. I hereby expressly consent to the personal jurisdiction of the state and federal courts located in California for any lawsuit filed there against me by the Company arising from or relating to this Agreement.

(b) Entire Agreement; Enforcement of Right. This Agreement sets forth the entire agreement and understanding between the Company and me relating to the subject matter herein and merges all prior discussion between us. No modification of or amendment to this Agreement, nor any waiver of any rights under this agreement, will be effective unless in writing signed by the party to be charged. The failure by either party to enforce any rights hereunder will not be construed as a waiver of any rights of such party. Any subsequent change or changes in my duties, salary or compensation will not affect the validity or scope of this Agreement.

(c) Severability. If one or more of the provisions in this Agreement are deemed void by law, then the remaining provisions will continue in full force and effect.

(d) Successors and Assigns. This will be binding upon my heirs, executors, administrators and other legal representatives and will be for the benefit of the Company, its successors, and its assigns.

Date: _____

Signature

Name of Employee (typed or printed)

Witness

EXHIBIT C

[Alcatel Letter Head]

Scot Reader
Alcatel Internetworking, Inc.
26801 West Agoura Road
Calabasas, California 91301

Via Federal Express

Eugene L. Hahm, Esq.
Coudert Brothers
530 Lytton Avenue, Suite 300
Palo Alto, California 94301-1541

Dear Mr. Hahm:

The purpose of this letter is to request that your clients Christopher Hoogenboom and John Wallner execute documents pertaining to certain patent applications on which they are listed as inventors. The applications claim inventions that your clients conceived while employed by Alcatel Internetworking, Inc., formerly Xylan Corp. Pursuant to their respective employment agreements, your clients are required to assign to Alcatel all rights in and to the inventions.

We are writing to you, as opposed to Messrs. Hoogenboom and Wallner directly, as we are aware that your firm represents them as litigation counsel.

If everything is in order, please have the inventors sign and date (exactly as the name is typed) the following enclosed documents relating to the enclosed patent applications:

- I. For U.S. Patent Application No. 09/718,046 entitled "Source Learning for Multi-Path Backplane Data Switch," please have Christopher Hoogenboom and John Wallner review and execute the enclosed Declaration and Power of Attorney and Assignment.
- II. For U.S. Patent Application No. 09/718,050, entitled "Flow Integrity for Path Transitioning Data Switch," please have Christopher Hoogenboom and John Wallner review and execute the enclosed Declaration and Power of Attorney and Assignment.
- III. For U.S. Patent Application No. 09/098,228, entitled "Digital Traffic Switch with Credit-Based Buffer Control," please have Christopher Hoogenboom execute the enclosed Assignment.

The text at the top of the Declaration and Power of Attorney forms should be carefully reviewed to ensure that the statements made are correct. The Assignments establish Alcatel as the owner of the respective applications and any and all patent rights in and to the respective inventions. The Assignment must be signed in front of a Notary

Public who is to enter the usual endorsement which must bear the same signature date as the Declaration and Power of Attorney.

Please note that the enclosed documents and the contents thereof are Alcatel confidential documents and are not to be disseminated to anyone other than the inventors and inventors counsel.

Please return the patent applications and executed Declaration and Power of Attorney and Assignment to us, and we will forward the executed documents to the Patent and Trademark Office.

Sincerely

Scot A. Reader

FLOW INTEGRITY FOR PATH TRANSITIONING DATA SWITCH

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of the filing date
5 of U.S. provisional patent application number 60/185,153,
filed February 18, 2000 and entitled "Flow Integrity for
Path Transitioning Data Switch," the contents of which are
hereby incorporated by reference. This application is
related to U.S. patent application No. 09/718,046 filed
10 November 20, 2000 and entitled "Source Learning for Multi-
Path Backplane Data Switch."

FIELD OF THE INVENTION

The present invention relates to devices for switching
15 data and, more particularly, to devices for switching data
while maintaining flow integrity.

BACKGROUND OF THE INVENTION

Data switches interconnect network devices residing in
20 different parts of a network. Such switches typically
include network interface modules having ports supporting
external network devices. In order to avoid having to
"flood" packets on ports of the switch which do not support
a destination device, such switches often dynamically learn
25 associations between network interface modules and the
devices they support. Such associations are often learned
through what is called a "source learning" process. In an
exemplary source learning process, source addresses in
packets are reviewed by the switch upon ingress and unknown
30 source addresses are submitted to a source learning
function resident on the switch. The source learning
function typically updates the switch such that future

packets destined to that source can be forwarded without unnecessary flooding.

While the source learning function has worked well in data switches, it typically gives rise to certain complications as the backplanes in such switches change from the single flow path variety to a multiple flow path variety. A conventional backplane used to interconnect network interface modules in a modular data switch has a common path for all packets between two network interface modules. All packets transmitted on the backplane for a particular flow typically have been propagated on a common path and the network interface modules have been responsible for individually reviewing the packets and making filtering decisions. However, due to demands for faster switching, there is an emerging trend favoring multiple flow path backplanes.

In a multiple flow path backplane architecture, packets destined for a destination device may be transmitted to a multicast fabric for queuing, replication, and forwarding to all network interface modules on a set of point-to-point paths before the destination device/network interface module association becomes known. On the other hand, packets destined for the destination device may be transmitted to a unicast fabric for queuing and forwarding to a single network interface module on a point-to-point path after the destination device/network interface module association becomes known.

Accordingly, in the case where destination device/network interface module associations are established through source learning, packets for a particular flow may be transmitted to the multicast fabric prior to learning the destination device's address. On the

other hand, packets for the flow may be transmitted to the unicast fabric once the destination device's address is learned. This dynamic transition from multicast to unicast fabrics, often made in connection with source learning,

5 gives rise to certain technical challenges. For instance, a packet transmitted to the unicast fabric at time t_2 may be forwarded to the network interface module associated with the destination device before a packet for the same flow transmitted to the multicast fabric at an earlier time t_1 .

10 This would result in the network interface module receiving and processing packets for the flow out-of-order, disrupting the integrity of the flow.

Therefore, there is a need for a multiple flow path data switch that is capable of maintaining flow integrity

15 during path transitions.

SUMMARY OF THE INVENTION

In one embodiment of the present invention, a path transitioning data switch is provided. The path

20 transitioning data switch has a plurality of switching modules and a backplane interconnecting the switching modules on a plurality of paths. A flow integrity of data units for a flow undergoing path transition is maintained by temporarily disabling one or more of the switching

25 modules from transmitting the data units for the flow to the backplane.

In another embodiment of the present invention, a path transitioning data switch having a plurality of switching modules is provided. A backplane interconnects the

30 switching modules on a plurality of paths. The data switch includes: means for transmitting a first data unit for a flow from a first switching module to a second switching

module on a first path; means for temporarily disabling the first switching module from transmitting data units for the flow; means for transitioning the flow from the first switching module to the second switching module to a second path; and means for transmitting a second data unit for the flow from the first switching module to the second switching module on the second path.

In yet another embodiment of the present invention, a path transitioning data switch having a plurality of switching modules is provided. A backplane interconnects the switching modules. The data switch includes: means for transmitting a first data unit having a first address as a destination address over the backplane; means for transmitting a second data unit having the first address as a source address over the backplane; and means in response to the second data unit for temporarily disabling transmission over the backplane of data units having the first address as a destination address.

In yet another embodiment of the present invention, a method of flow path transitioning in a data communication switch is provided. The data communication switch has a plurality of flow paths and a plurality of network interface modules. A first packet having an unknown source address is received from a source device coupled to a first network interface module. The first packet is transmitted over a first flow path to one or more network interface modules. Source learning of the source address is performed. The first network interface module is disabled from transmitting packets from the source device to other network interface modules while source learning of the source address is being performed.

In yet another embodiment of the present invention, a data communication switch having a backplane is provided. A plurality of network interface modules are interconnected over the backplane. Each network interface module includes: an access controller having a port for receiving a plurality of packets; a switching controller coupled to the access controller for receiving the packets from the access controller, and for processing the packets for routing; and a fabric controller coupled to the switching controller for receiving the packets from the switching controller, and for processing the packets for transmitting over the backplane to one or more other network interface modules. The fabric controller receives packets from the other network interface modules and provides them to the switching controller, and the switching controller provides the packets from the other network interface modules to the access controller for transmitting out of the port.

In yet another embodiment of the present invention, a method of maintaining flow integrity in a data communication switch is provided. The data communication switch has a plurality of network interface modules coupled over a backplane. Each network interface module has a plurality of ports. A packet of a flow is received on a first network interface module. The packet is discarded if both source and destination addresses are found. It is determined as to whether only the source address is found unless both the source address and the destination address are found.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects of the invention can be understood by reference to the following detailed description, taken in conjunction with the accompanying drawings, which are briefly described below.

Figure 1 is a block diagram illustrating a data communication switch with transitioning flow path;

Figure 2 is a block diagram illustrating in more detail a representative network interface module within the switch of Figure 1;

Figure 3 is a flow diagram illustrating a source and destination indexing protocol in the switch according to Figure 1;

Figure 4 is a flow diagram illustrating a path transitioning protocol in the switch according to Figure 1;

Figure 5 is a flow diagram illustrating the unicast receive protocol in the switch according to Figure 1;

Figure 6 is a flow diagram illustrating the multicast receive protocol in the switch according to Figure 1;

Figure 7 is a flow diagram illustrating the management interface protocol in the switch according to Figure 1; and

Figure 8 is a flow diagram illustrating instructions provided to the network interface modules of Figure 7 in further detail.

DETAILED DESCRIPTION

Figure 1 illustrates a switch 100 including network interface modules 110, 120 and 130. The switch may also be referred to as a data switch, as a data communication switch, or as other designations used by those skilled in the art. The network interface modules may also be referred to as switching modules, or as other designations

used by those skilled in the art. The network interface modules 110, 120 and 130 are operatively coupled to LANs 111, 121, and 131, respectively. The switch 100 also includes a management interface module 140, a unicast
5 fabric 150, a multicast interface 160 and a multicast fabric 170. The network interface modules 110, 120, and 130 communicate with the unicast fabric 150 via bi-directional signals 112, 122, and 132, and preferably serve to associate routing information with addresses, cellify
10 data packets, and perform other look-up functions.

The unicast fabric 150 preferably is used as the routing link when the source and destination addresses are known. The multicast interface 160 is operatively coupled to the unicast fabric 150 and to the multicast fabric 170.
15 The multicast interface 160 preferably receives an input signal from the unicast fabric 150 and provides an output signal to the multicast fabric 170. The multicast interface 160 may also receive other signals and provide them to the multicast fabric 170. The output signal may
20 include the input signal, the other signals, or a composite of the input signals and some or all of the other signals. The multicast fabric 170 preferably provides the input signal to the network interface modules for source learning when either the source address or the destination address
25 is unknown.

The multicast fabric 170 is operatively coupled to the multicast interface 160, to the management interface module 140, and to the network interface modules 110, 120, and 130. The multicast fabric 170 preferably is used as the
30 routing link when either the source address or the destination address is not known, or when both the source and destination addresses are not known. The management

interface module 140 is operatively coupled to the multicast fabric 170 and to the network interface modules 110, 120, and 130. The management interface module 140 preferably is used to discard packets and update addresses
5 as required. It will be appreciated that the multicast interface 160 and the management interface module 140 may share the same logic block.

Figure 2 is a block diagram of a network interface module 200, which is representative of all network
10 interface modules in one embodiment of the present invention. The network interface module 200 includes an access controller 210, a switching controller 220, and a fabric controller 230. The access controller 210 preferably includes a simple physical layer device (PHY
15 device), which is coupled to LANs 205. The access controller 210 preferably receives packets off LANs 205 and transmits them to the switching controller 220. The access controller 210 preferably also receives packets from the switching controller 220 and transmits them to the LANs
20 205.

The switching controller 220 is coupled to modules for facilitating flow integrity, including a switching data module 222 and a content addressable memory (CAM) 221. The switching controller 220 preferably receives packets from
25 the access controller 210, and performs a local CAM lookup using the CAM 221 to retrieve CAM data. The switching controller preferably also performs a switching data lookup by providing CAM data to the switching data module 222 to retrieve switching data. The switching data may also be
30 referred to as forwarding data. Further, the switching controller uses the CAM data and the switching data to process the packets, subjects the packets to flow

integrity, and forwards the packets to the fabric controller 230.

The switching controller 220 is also coupled to a switching control update signal 142, which enables the switching controller 220 to update the CAM 221 with further addresses. The fabric controller 230 is coupled to modules for facilitating flow integrity, including a fabric data module 233, a pseudo CAM 232, and a VLAN finder module 231. The fabric controller 230 preferably receives packets from the switching controller 220, and preferably performs a local CAM lookup with either the pseudo CAM 232 or the VLAN finder module 231. The fabric controller preferably also subjects the packets to ingress flow integrity such that the packets are converted into cells, and then preferably forwards those cells out of the network interface module 200.

The fabric controller 230 is also coupled to a fabric control update signal 141 which preferably enables the fabric controller 230 to discard packets for a period of time, which, for example, may have been predetermined. This discard function preferably allows time for the switching controller 220 to update the CAM 221 with a new address. Returning now to Figure 1, the discard function preferably also allows packets from the multicast fabric 170 to arrive at their destination before packets from the unicast fabric 150 are allowed to propagate to the same destination.

In Figures 3 through 7, flow integrity is described in reference to flow diagrams.

Source and Destination Conditions:

One embodiment of the present invention is directed to a novel flow integrity technique using multiple flow path backplanes while performing source learning. In a multiple flow path backplane architecture, packets with either an unknown destination address or an unknown source address are transmitted to a multicast fabric for queuing, replication, and forwarding. On the other hand, packets with known destination and source addresses are transmitted to a unicast fabric for queuing and forwarding to a single network interface module on a point-to-point path. Consequently, the switch in this embodiment preferably is capable of processing and routing packets, wherein each packet presents one of the following four properties to the network interface modules: (1) unknown destination and unknown source; (2) unknown destination and known source; (3) known destination and unknown source; and (4) known destination and known source. In the case where the source and destination addresses are known, no source learning typically takes place and therefore flow integrity is not an issue.

Unknown Destination and Unknown Source:

Referring to Figure 3, a packet is received on a physical port of a network interface module (310), which may be similar to the network interface module 200 of Figure 2. Upon receiving this packet, a look-up operation preferably is performed to find source and destination addresses (320). Unless both the source and destination addresses are found (330), another query preferably is performed, as is the case here, to see if just the source address is found (340). If the source address is not found (340), as is the case here, a source virtual port number

(SVPN) preferably is applied to the packet, and a source CAM index (SCI) preferably is set as invalid (350).

Referring to Figure 4, the packet preferably is operated on to determine its virtual local area network (VLAN) identification (ID) from the SCI or SVPN (405). In this case, the VLAN ID preferably is determined from the SVPN since the SCI has been set as invalid. Next, an SCI-VLAN entry is looked for in the VLAN finder module using the VLAN ID to determine whether a discard indicator is set in the SCI-VLAN entry (410). The setting of the discard indicator in the SCI-VLAN entry preferably indicates that the packet is from the network device whose address is undergoing source learning. If the discard indicator is set (415), the packet preferably is discarded (475), and packets from the same network device preferably are discarded until the discard indicator is no longer set.

Under conditions here, the discard indicator is not set and, as a result, the packet preferably is not discarded. A pseudo CAM look-up key preferably is generated from the destination address and the VLAN ID (420). The pseudo CAM look-up key preferably is used to determine a destination CAM index (DCI) and a fabric data index (FDI) (425). Since the destination address is unknown in this case, a default destination CAM index (DCI) preferably is used.

The pseudo CAM look-up key preferably is also used to retrieve a pseudo CAM entry. If a discard indicator is set in the retrieved pseudo CAM entry (430), the packet preferably is discarded (475). The setting of the discard indicator in the pseudo CAM entry preferably indicates that the packet is destined to the network device whose address is undergoing source learning. Thus, as long as the

discard indicator is set, all packets destined to the same network device preferably are discarded. The discard indicator is not set in the pseudo CAM entry in this case, and the packet preferably is not discarded. The FDI
5 preferably is used to determine a destination interface ID and a multicast group ID (435).

Following this, the DCI preferably is applied to the packet (445). Since the destination address is unknown in this case, the DCI preferably has been set as default,
10 which preferably instructs a flood of the network interface modules. After the DCI is applied to the packets, the packets preferably are segmented into cells (450). The cells preferably are then given a source interface ID, the destination interface ID, and the multicast group ID (455).
15 In this case, before the cells are transmitted (470) over the multicast fabric, the destination interface ID preferably is set to the multicast group ID and a source learning indicator preferably is set in the cells (465) since the SCI is invalid(460).

20 Referring to Figure 6, on receipt of the cells from the multicast fabric (610), the network interface module preferably performs a look-up operation of the multicast group ID (620) to determine whether or not to transmit the cells, i.e., the packet, out of its ports. If the
25 multicast group ID is found (630), the cells preferably are reassembled (520) into the packet as illustrated in Figure 5. If the DCI indicates flood (530), as is the case here, the packet preferably is transmitted out of all ports (532). For example, the packets preferably flood all ports
30 of the network interface module if the DCI is the default DCI.

When the network interface modules receive the transmitted cells (610) as illustrated in Figure 6, the management interface module preferably also receives the transmitted cells (710) as illustrated in Figure 7. Once
5 the cells have been received by the management interface module, an operation preferably is performed to determine if the source learning indicator is set (720). In this case, the source learning indicator has been set and, as such, the management interface module preferably instructs
10 the network interface modules to perform a number of functions (750) including but not limited to the following.

The instructions to perform these functions illustrated in Figure 8 are described in reference to the network interface module 200 of Figure 2 as the network
15 interface module that interfaces with the network device whose address is undergoing source learning. However, the description of the functions are broadly applicable to any network interface module that interfaces with such network device, such as, for example, the network interface module
20 110, 120 or 130 of Figure 1.

First, the switching controller 220 coupled to the network device whose address is undergoing source learning preferably is instructed by the switching control update signal 142 to add a source address of the network device to
25 the CAM 221 (800). Second, the fabric controller 230 preferably is instructed by the fabric control update signal 141 to add a SCI-VLAN entry including a discard indicator to the VLAN Finder Module 231 (810). The discard indicator preferably causes the fabric controller 230 to
30 discard packets from the network device whose address is undergoing source learning for a period, which, for example, may have been predetermined. In other embodiments,

the period for discarding packets may be dynamically determined. Third and fourth, the management interface module preferably instructs the fabric controller 230 using the fabric control update signal 141 to add a pseudo CAM entry to the pseudo CAM 232 (820) and a fabric data entry (830) to the fabric data module 233, respectively.

Fifth, the management interface module preferably instructs the network interface modules that are not coupled to the network device whose address is undergoing source learning to add a pseudo CAM entry including a discard indicator to their respective pseudo CAMs using their respective fabric control update signals (840). The discard indicator preferably causes the respective fabric controllers of these network interface modules to discard packets destined to the network device whose address is undergoing source learning. Sixth, the management interface module preferably instructs the respective fabric controllers of the network interface modules that are not interfacing with the network device whose address is undergoing source learning to add a fabric data entry to their respective fabric data modules (850).

Once the discard conditions are indicated in the network interface modules for packets associated with a particular network device, the discard conditions preferably are removed after an interval passes to ensure that all the packets associated with that network device clear the multicast fabric. Here, the interval, for example, may have been predetermined. After the discard conditions are removed, the packets associated with that network device preferably are allowed to flow over the unicast fabric since both the destination and source addresses are now known to the switch.

Unknown Destination and Known Source:

Referring to Figure 3, a packet is received on a physical port of a network interface module (310), which
5 may be similar to the network interface module 200 of Figure 2. Upon receiving this packet, a look-up operation preferably is performed to find source and destination addresses (320). Unless both the source and destination addresses are found (330), another query preferably is
10 performed, as is the case here, to see if just the source address is found (340). If the source address is found (340), as is the case here, the packet preferably is outputted (400) with the source CAM index (SCI) applied (342).

15 Referring to Figure 4, the packet (400) preferably is operated on to determine its virtual local area network (VLAN) identification (ID) from the SCI or SVPN (405). In this case, the VLAN ID preferably is determined from the SCI since the source address is known. Next, an SCI-VLAN
20 entry is looked for in the VLAN finder module using the VLAN ID to determine whether a discard indicator is set in the SCI-VLAN entry (410). If the discard indicator is set (415), the packet preferably is discarded (475), and packets from the same network device as the packet
25 preferably are discarded until the discard indicator is no longer set.

Under conditions here, the discard indicator is not set and, as a result, the packet preferably is not discarded. A pseudo CAM look-up key preferably is
30 generated from the destination address and the VLAN ID (420). The pseudo CAM look-up key preferably is used to determine a destination CAM index (DCI) and a fabric data

index (FDI) (425). Since the destination address is unknown in this case, a default destination CAM index (DCI) preferably is used.

The pseudo CAM look-up key preferably is also used to retrieve a pseudo CAM entry. If a discard indicator is set in the retrieved pseudo CAM entry (430), the packet preferably is discarded (475). As long as the discard indicator is set, all packets destined to the same network device as the packet preferably are discarded. The discard indicator is not set in the pseudo CAM entry in this case, and the packet preferably is not discarded. The FDI preferably is used to determine a destination interface ID and a multicast group ID (435).

Following this, the DCI preferably is applied to the packet (445). Since the destination address is unknown in this case, the DCI preferably has been set as default, which preferably instructs a flood of the network interface modules. After the DCI is applied to the packets, the packets preferably are segmented into cells (450). Prior to being transmitted (470), the cells preferably are given a source interface ID, the destination interface ID, and the multicast group ID (455).

Referring to Figure 6, on receipt of the cells from the multicast fabric (610), the network interface module preferably performs a look-up operation of the multicast group ID (620) to determine whether or not to transmit the cells, i.e., the packet, out of its ports. If the multicast group ID is found (630), the cells preferably are reassembled (520) into the packet as illustrated in Figure 5. If the DCI indicates flood (530), as is the case here, the packet preferably is transmitted out of all ports (532).

When the network interface modules receive the transmitted cells (610) as illustrated in Figure 6, the management interface module preferably also receives the transmitted cells (710) as illustrated in Figure 7. Once the cells have been received by the management interface module, an operation preferably is performed to determine if the source learning indicator is set (720). In this case, the source learning indicator has not been set and, as such, the management interface module preferably queries to see if the multicast group ID is found (730). In this case, the multicast group ID is found and the packet preferably is reassembled and processed (520) for receipt as illustrated in Figure 5. Since the default DCI instructs the network interface module to flood (530) in this case, the packet is transmitted out of all ports (532) of the network interface module.

Known Destination and Unknown Source:

Referring to Figure 3, a packet is received on a physical port of a network interface module (310), which may be similar to the network interface module 200 of Figure 2. Upon receiving this packet, a look-up operation preferably is performed to find source and destination addresses (320). Unless both the source and destination addresses are found (330), another query preferably is performed, as is the case here, to see if just the source address is found (340). If the source address is not found (340), as is the case here, a source virtual port number (SVPN) preferably is applied to the packet, and a source CAM index (SCI) preferably is set as invalid (350).

Referring to Figure 4, the packet preferably is operated on to determine its virtual local area network

(VLAN) identification (ID) from the SCI or SVPN (405). In this case, the VLAN ID preferably is determined from the SVPN since the SCI has been set as invalid. Next, an SCI-VLAN entry is looked for in the VLAN finder module using
5 the VLAN ID to determine whether a discard indicator is set in the SCI-VLAN entry (410). If the discard indicator is set (415), the packet preferably is discarded (475), and packets from the same network device as the packet preferably are discarded until the discard indicator is no
10 longer set.

Under conditions here, the discard indicator is not set and, as a result, the packet preferably is not discarded. A pseudo CAM look-up key preferably is generated from the destination address and the VLAN ID
15 (420). The pseudo CAM look-up key preferably is used to determine a destination CAM index (DCI) and a fabric data index (FDI) (425). Since the destination address is known in this case, a default destination CAM index (DCI) preferably is not used.

20 The pseudo CAM look-up key preferably is also used to retrieve a pseudo CAM entry. If a discard indicator is set in the retrieved pseudo CAM entry (430), the packet preferably is discarded (475). As long as the discard indicator is set, all packets destined to the same network
25 device as the packet preferably are discarded. The discard indicator is not set in the pseudo CAM entry in this case, and, the FDI preferably is used to determine a destination interface ID and a multicast group ID (435).

Following this, the DCI preferably is applied to the
30 packet (445) and the packets preferably are segmented into cells (450). The cells preferably are then given a source interface ID, the destination interface ID, and the

multicast group ID (455). In this case, before the cells are transmitted (470) over the multicast fabric, the destination interface ID preferably is set to the multicast group ID and a source learning indicator preferably is set
5 in the cells (465) since the SCI is invalid.

Referring to Figure 6, on receipt of the cells from the multicast fabric (610), the network interface module preferably performs a look-up operation of the multicast group ID (620) to determine whether or not to transmit the
10 cells, i.e., the packet, out of its ports. If the multicast group ID is found (630), the cells preferably are reassembled into the packet (520) as illustrated in Figure 5. If the DCI indicates flood (530), the packet preferably is transmitted of all ports (532). In this case, since the
15 destination address is known, the default DCI preferably is not used, and thus the packet preferably is not transmitted out of all ports.

Instead, a forwarding port ID preferably is determined from the DCI (540). Then the DCI and the SVPN preferably
20 are used to determine whether or not the source and destination devices of the packet share a virtual LAN (VLAN) (550). If a VLAN is shared (560), the packet preferably is transmitted on the forwarding port identified by the forwarding port ID (570). If, however, the VLAN is
25 not shared by the source and destination devices, the packet preferably is discarded (562).

When the network interface modules receive the transmitted cells (610) as illustrated in Figure 6, the management interface module preferably also receives the
30 transmitted cells (710) as illustrated in Figure 7. Once the cells have been received by the management interface module, an operation preferably is performed to determine

if the source learning indicator is set (720). In this case, the source learning indicator has been set and, as such, the management interface module preferably instructs the network interface modules to perform a number of functions (750) including but not limited to the following.

The instructions to perform these functions illustrated in Figure 8 are described in reference to the network interface module 200 of Figure 2 as the network interface module that interfaces with the network device whose address is undergoing source learning. However, the description of the functions are broadly applicable to any network interface module that interfaces with such network device such as, for example, the network interface module 110, 120 or 130 of Figure 1.

First, the switching controller 220 coupled to the network device whose address is undergoing source learning preferably is instructed by the switching control update signal 142 to add a source address of the network device to the CAM 221 (800). Second, the fabric controller 230 preferably is instructed by the fabric control update signal 141 to add a SCI-VLAN entry including a discard indicator to the VLAN Finder Module 231 (810). The discard indicator preferably causes the fabric controller 230 to discard packets from the network device whose address is undergoing source learning for a period, which, for example, may have been predetermined. In Other embodiments, the period for discarding packets may be dynamically determined. Third and fourth, the management interface module preferably instructs the fabric controller 230 using the fabric control update signal 141 to add a pseudo CAM entry to the pseudo CAM 232 (820) and a fabric data entry (830) to the fabric data module 233, respectively.

Fifth, the management interface module preferably instructs the network interface modules that are not coupled to the network device whose address is undergoing source learning to add a pseudo CAM entry including a discard indicator to their respective pseudo CAMs using their respective fabric control update signals (840). The discard indicator preferably causes the respective fabric controllers of these network interface modules to discard packets destined to the network device whose address is undergoing source learning. Sixth, the management interface module preferably instructs the respective fabric controllers of the network interface modules that are not interfacing with the network device whose address is undergoing source learning to add a fabric data entry to their respective fabric data modules (850).

Once the discard conditions are indicated in the network interface modules for packets associated with a particular network device, the discard conditions preferably are removed after an interval passes to ensure that all packets associated with that network device clear the multicast fabric. Here, the interval, for example, may have been predetermined. After the discard conditions are removed, the packets associated with that network device preferably are allowed to flow over the unicast fabric since both the destination and source addresses are now known to the switch.

Known Destination and Known Source:

Referring to Figure 3, a packet is received on a physical port of a network interface module (310), which may be similar to the network interface module 200 of Figure 2. Upon receiving this packet, a look-up operation

preferably is performed to find source and destination addresses (320).

There are two different cases here, depending on whether the source and destination devices are local to the same network interface module or not. If the destination and the source devices are coupled to the same network interface module, the source and destination addresses are found (330) and the packet is discarded (332) from being sent to other network interface modules over the fabrics.

If the source and destination devices are not coupled to the same network interface module, the source address is typically locally known (340) but the destination address may not be locally known. A source CAM index (SCI) preferably is applied (342) to the packet.

Referring to Figure 4, the packet preferably is operated on to determine its virtual local area network (VLAN) identification (ID) from the SCI or SVPN (405). In this case, the VLAN ID preferably is determined from the SCI since the source address is known. Next, an SCI-VLAN entry is looked for in the VLAN finder module using the VLAN ID to determine whether a discard indicator is set in the SCI-VLAN entry (410). The setting of the discard indicator in the SCI-VLAN entry preferably indicates that the packet is from the network device whose address is undergoing source learning.

Since both the source and destination addresses are known in this case, no source learning associated with the packet preferably takes place, and thus the discard indicator preferably is not set. Then a pseudo CAM look-up key preferably is generated from the destination address and the VLAN ID (420). Next, the DCI and FDI preferably are determined from the pseudo CAM look-up key (425). Since

the destination address is known in this case, a default destination CAM index (DCI) preferably is not used.

The pseudo CAM look-up key preferably is also used to retrieve a pseudo CAM entry. If a discard indicator is set
5 in the retrieved pseudo CAM entry (430), the packet preferably is discarded (475). The setting of the discard indicator in the pseudo CAM entry preferably indicates that the packet is destined to the network device whose address is undergoing source learning. In this case, both the
10 source and destination address are known for the packet, and thus no source learning associated with the packet preferably takes place, and the discard indicator preferably is not set.

Next, the FDI preferably is used to determine a
15 destination interface ID and a multicast group ID (435). Thereafter, the DCI preferably is applied to the packet (445) and then the packets are segmented into cells (450). Prior to being transmitted (470), the cells preferably are given a source interface ID, the destination interface ID,
20 and the multicast group ID (455). Since both the source and the destination addresses are known, the multicast group ID in this case preferably contains a single destination. Therefore, the cells preferably are not forwarded to the multicast fabric but preferably are
25 forwarded to the unicast fabric.

Referring to Figure 5, on receipt of the cells from the unicast fabric (510), the cells preferably are reassembled into the packet (520). Next, a test preferably is performed to determine if the DCI has been set to flood
30 (530). In this, the DCI preferably has not been set to flood since both the source and destination addresses are

known and the unicast fabric has been used. Therefore, the packet is not transmitted out of all ports.

Instead, a forwarding port ID preferably is determined from the DCI (540). Then the DCI and the SCI preferably
5 are used to determine whether or not the source and destination devices of the packet share a virtual LAN (VLAN) (550). If a VLAN is shared (560), the packet preferably is transmitted on the forwarding port identified by the forwarding port ID (570). If, however, the VLAN is
10 not shared by the source and destination devices, the packet preferably is discarded (562).

It will be appreciated by those of ordinary skill in the art that the invention can be embodied in other specific forms without departing from the spirit or
15 essential character hereof. The present description is therefore considered in all respects to be illustrative and not restrictive. The scope of the invention is indicated by the appended claims, and all changes that come within the meaning and range of equivalents thereof are intended
20 to be embraced therein.

We claim:

1. A path transitioning data switch comprising:
a plurality of switching modules; and
a backplane interconnecting the switching modules
5 on a plurality of paths,
wherein flow integrity of data units for a flow
undergoing path transition is maintained by temporarily
disabling one or more of the switching modules from
transmitting the data units for the flow to the backplane.

10

2. The path transitioning data switch of claim 1
wherein the switching modules are disabled from
transmitting the data units for the flow when the path
transition is commenced, and are enabled after the path
15 transition has been completed and an interval has passed to
ensure that all data units for the flow transmitted to the
backplane prior to disabling the switching modules have
cleared the backplane.

20

3. The path transitioning data switch of claim 1
wherein the backplane includes a multicast fabric and a
unicast fabric, and wherein the path transition is made
from the multicast fabric to the unicast fabric upon source
learning an address of a network device that provides the
25 flow.

25

4. The path transitioning data switch of claim 3
wherein all switching modules are temporarily disabled from
transmitting data units having the address undergoing
30 source learning as a destination address to the backplane.

5 5. The path transitioning data switch of claim 4 wherein the switching module coupled to the network device whose address is undergoing source learning is temporarily disabled from transmitting data units having the address as a source address to the backplane.

10 6. The path transitioning data switch of claim 5 wherein the temporarily disabled switching modules are enabled to transmit to the backplane after source learning has been completed.

15 7. The path transitioning data switch of claim 6 wherein the temporarily disabled switching modules are enabled after an interval has passed to ensure that all data units having the source learned address as the source address or the destination address and transmitted to the multicast fabric have cleared the backplane.

20 8. The path transitioning data switch of claim 7 wherein said interval is predetermined.

25 9. A path transitioning data switch having a plurality of switching modules and a backplane interconnecting the switching modules on a plurality of paths, the data switch comprising:

 means for transmitting a first data unit for a flow from a first switching module to a second switching module on a first path;

30 means for temporarily disabling the first switching module from transmitting data units for the flow;

means for transitioning the flow from the first switching module to the second switching module to a second path; and

means for transmitting a second data unit for the
5 flow from the first switching module to the second switching module on the second path.

10. The path transitioning data switch of claim 9 wherein the first switching module includes the means for
10 transmitting a first data unit and the means for transmitting a second data unit.

11. The path transitioning data switch of claim 9 wherein the means for transmitting a first data unit
15 transmits the data units for the flow on the first path before the flow is transitioned from the first path to the second path, and wherein the means for transmitting a second data unit transmits the data units for the flow on the second path after the flow has been transitioned from
20 the first path to the second path.

12. The path transitioning data switch of claim 11 wherein the first path includes a multicast fabric, the second path includes a unicast fabric, and the means for
25 transitioning the flow transitions the flow from the multicast fabric to the unicast fabric after an address of a network device that provides the flow undergoes source learning.

30 13. The path transitioning data switch of claim 12 wherein the first switching module is coupled to the network device that provides the flow, and the means for

temporarily disabling the first switching module prevents the first switching module from transmitting the flow from the network device to the backplane while the address of the network device is undergoing source learning.

5

14. A path transitioning data switch of claim 12 wherein the first switching module is coupled to the network device that provides the flow, and the means for temporarily disabling the first switching module prevents
10 the first switching module from transmitting the flow from the network device to the backplane for an interval to ensure that all data units of the flow have cleared the multicast fabric.

15 15. The path transitioning data switch of claim 14 wherein said interval is predetermined.

16. The path transitioning data switch of claim 13 wherein the means for temporarily disabling the first
20 switching module includes means for discarding data units and means for instructing the means for discarding data units to discard the data units for the flow.

17. The path transitioning data switch of claim 16
25 wherein the means for temporarily disabling the first switching module further includes means for instructing the means for discarding data units to stop discarding the data units for the flow.

30 18. The path transitioning data switch of claim 14 wherein the means for transmitting a second data unit

transmits the second data unit on the unicast fabric after the source learning has been completed.

19. A path transitioning data switch having a
5 plurality of switching modules and a backplane interconnecting the switching modules, the data switch comprising:

means for transmitting a first data unit having a first address as a destination address over the backplane;

10 means for transmitting a second data unit having the first address as a source address over the backplane; and

means in response to the second data unit for temporarily disabling transmission over the backplane of
15 data units having the first address as a destination address.

20. A path transitioning data switch of claim 19 wherein the first address is associated with a network
20 device coupled to a first switching module, and the data switch performs source learning of the first address after receiving the second data unit from the network device.

21. A path transitioning data switch of claim 20
25 wherein the means for temporarily disabling transmission prevents transmission of the data units having the first address as the destination address while the first address undergoes source learning.

30 22. The path transitioning data switch of claim 21 wherein the means for temporarily disabling transmission includes means for discarding data units and means for

instructing the means for discarding data units to discard the data units having the first address as the destination address.

5 23. The path transitioning data switch of claim 22 wherein the means for temporarily disabling transmission further includes means for instructing the means for discarding data units to stop discarding the data units having the first address as the destination address.

10

24. A path transitioning data switch of claim 20 wherein the means for temporarily disabling transmission prevents transmission of the data units from the network device to the backplane for an interval to ensure that all
15 data units having the first address as the destination address have cleared the backplane.

25. The path transitioning data switch of claim 24 wherein said interval is predetermined.

20

26. A method of flow path transitioning in a data communication switch having a plurality of flow paths and a plurality of network interface modules, the method comprising the steps of:

25 receiving a first packet having an unknown source address from a source device coupled to a first network interface module;

transmitting the first packet over a first flow path to one or more network interface modules; and

30 performing source learning of the source address, wherein the first network interface module is

disabled from transmitting packets from the source device to other network interface modules while source learning of the source address is being performed.

5 27. The method of flow path transitioning of claim 26 wherein other network devices are disabled from transmitting any packet having the source address undergoing source learning as a destination address over the flow paths until source learning has been completed.

10

28. The method of flow path transitioning of claim 26, the method further comprising the step of:

transmitting a second packet from the source device over a second flow path to a second network switching module after the source learning has been completed,

wherein the first flow path includes a multicast fabric and the second flow path includes a unicast fabric.

20 29. A data communication switch having a backplane and a plurality of network interface modules interconnected over the backplane, each network interface module comprising:

an access controller having a port for receiving a plurality of packets;

a switching controller coupled to the access controller for receiving the packets from the access controller, and for processing the packets for routing; and

a fabric controller coupled to the switching controller for receiving the packets from the switching controller, and for processing the packets for transmitting

over the backplane to one or more other network interface modules,

wherein the fabric controller receives packets from the other network interface modules and provides them to the switching controller, and the switching controller provides the packets from the other network interface modules to the access controller for transmitting out of the port.

30. The data communication switch of claim 29, wherein the backplane includes a unicast fabric, a multicast fabric and a management interface module, and wherein the fabric controller is used to transmit the packets over either the unicast fabric or the multicast fabric.

31. The data communication switch of claim 30 wherein the management interface module provides a fabric control update signal to the fabric controller with an instruction to discard packets for a flow undergoing a path transition from the multicast fabric to the unicast fabric, and wherein the fabric controller discards such packets.

32. The data communication switch of claim 30 wherein the management interface module provides a fabric control update signal to the fabric controller with an instruction to discard packets having a source address of a flow undergoing a path transition from the multicast fabric to the unicast fabric as a destination address, and wherein the fabric controller discards such packets.

33. A method of maintaining flow integrity in a data communication switch having a plurality of network interface modules coupled over a backplane, each network interface module having a plurality of ports, the method
5 comprising:

receiving a packet of a flow on a first network interface module;

discarding the packet if both source and destination addresses are found; and

10 determining if only the source address is found unless both the source address and the destination address are found.

34. The method of claim 33 wherein the backplane
15 includes a unicast fabric and a multicast fabric, the method further comprising the steps of:

transitioning a flow path of the flow from the multicast fabric to the unicast fabric.

20 35. The method of claim 34 wherein the step of transitioning a flow path includes the steps of: performing source learning of the flow if the source address has not been found;

disabling the first network interface module from
25 transmitting any packet of the flow to the backplane during source learning; and

disabling all other network interface modules from transmitting any packet having the source address as a destination address to the backplane during source
30 learning.


36. The method of claim 35 wherein the step of disabling the first network interface module includes the step of providing a discard indicator to the first network interface module, wherein the first network interface
5 module discards all packets of the flow in response to the discard indicator.

37. The method of claim 35 wherein the step of disabling all other network interface modules includes the
10 step of providing a discard indicator to all other network interface modules, wherein all other network interface modules discard all packets having the source address as a destination address.

ABSTRACT OF THE DISCLOSURE

A data switch includes multiple switching modules interconnected over a backplane. The data switch maintains flow integrity while path transitioning. The flow integrity
5 can be maintained by temporarily disabling one or more of the switching modules from transmitting data units over the backplane for a flow undergoing a path transition. The disable condition is imposed when the path transition is commenced, and is lifted after the path transition is
10 completed and an interval has passed ensuring that all data units for the flow transmitted to the backplane prior to imposing the disable condition have cleared the backplane. The backplane includes a multicast fabric and a unicast fabric, wherein flow path transitions are made from the
15 multicast fabric to the unicast fabric upon source learning. All switching modules are temporarily disabled from transmitting data units having as a destination address an address undergoing source learning to the backplane. Moreover, the switching module supporting the
20 device whose address is undergoing source learning is temporarily disabled from transmitting data units having the address as a source address to the backplane. The disable conditions are lifted after source learning has been completed and an interval has passed ensuring that all
25 data units having the source learned address as a source or destination address and transmitted to the multicast fabric have cleared the backplane.

FIGURE 1



ALCATEL CONFIDENTIAL


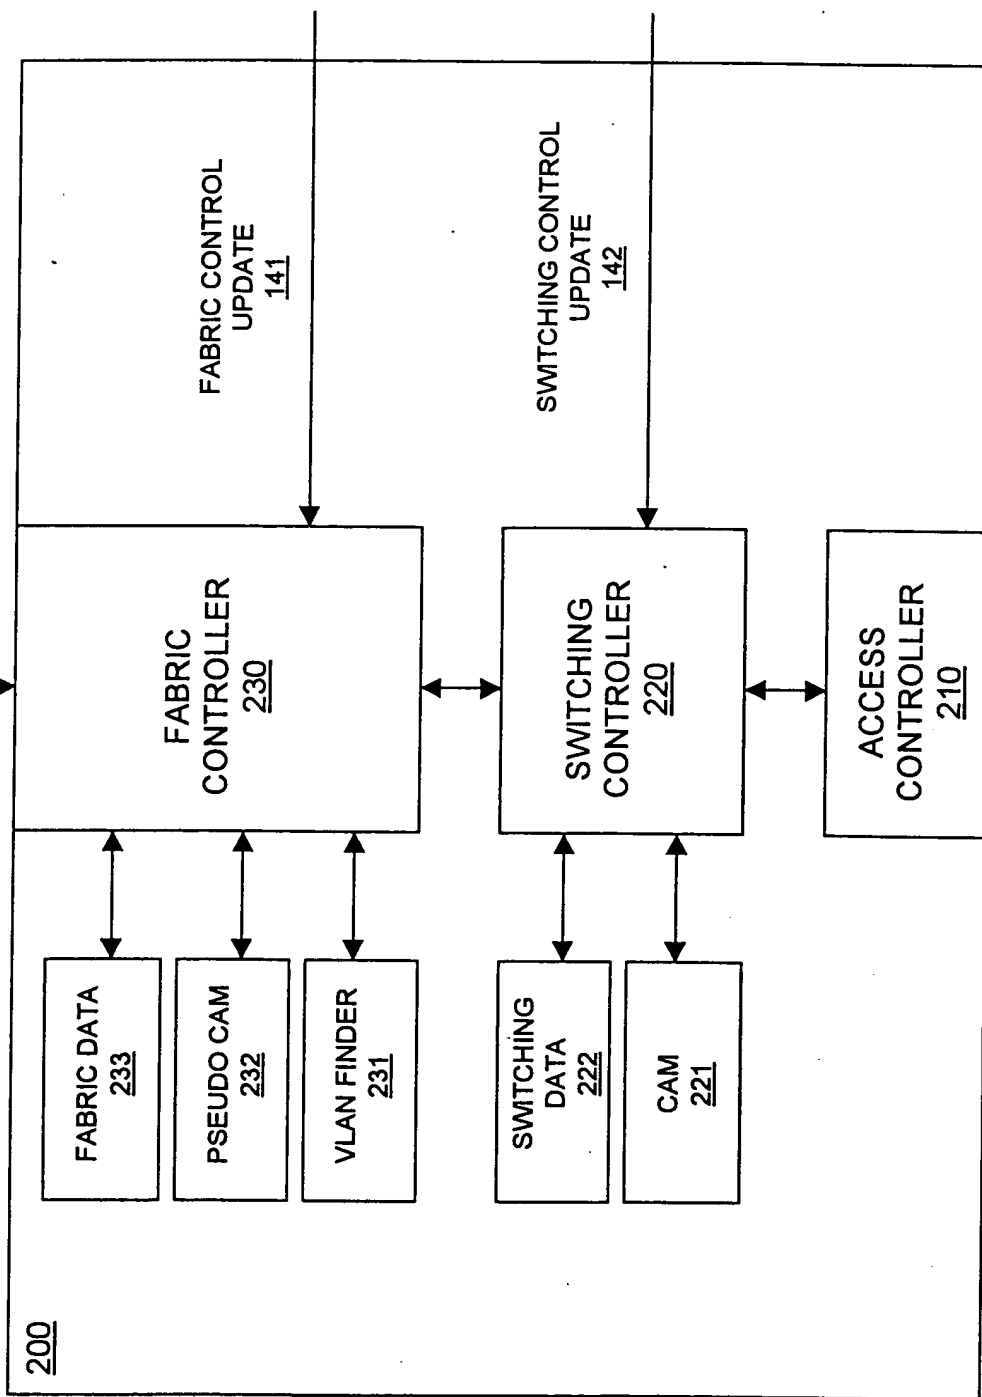



FIGURE 2

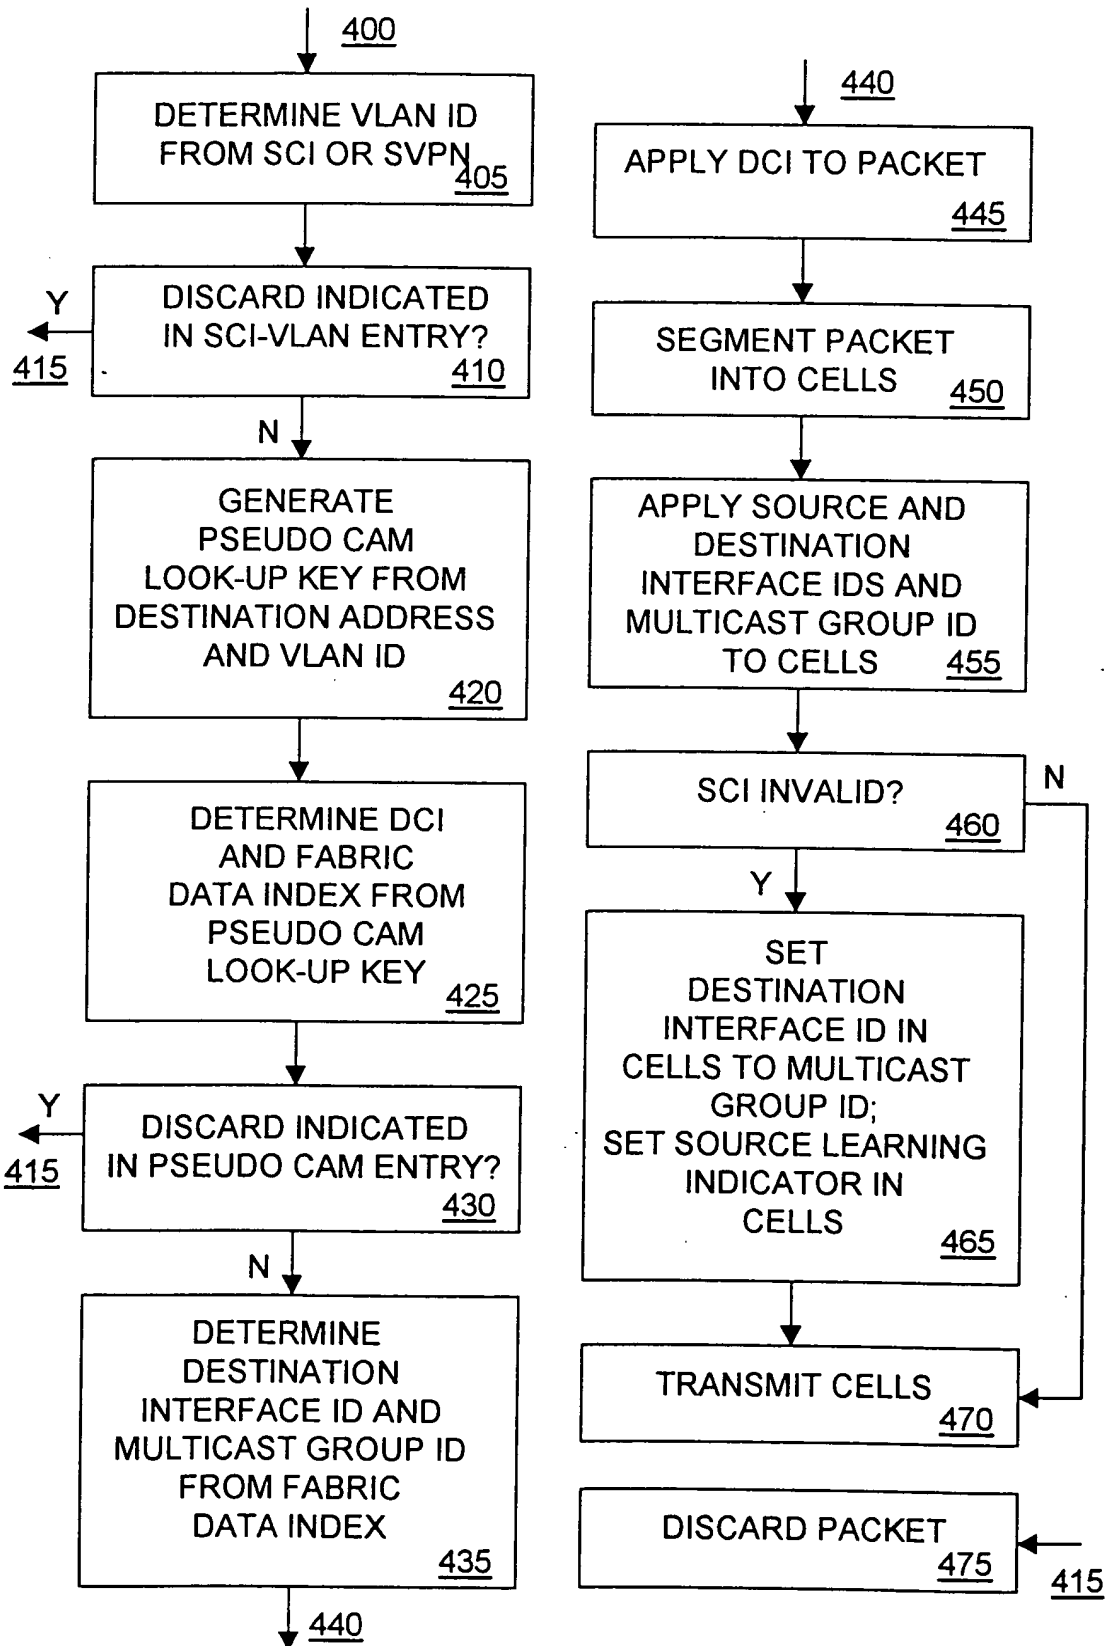


ALCATEL CONFIDENTIAL

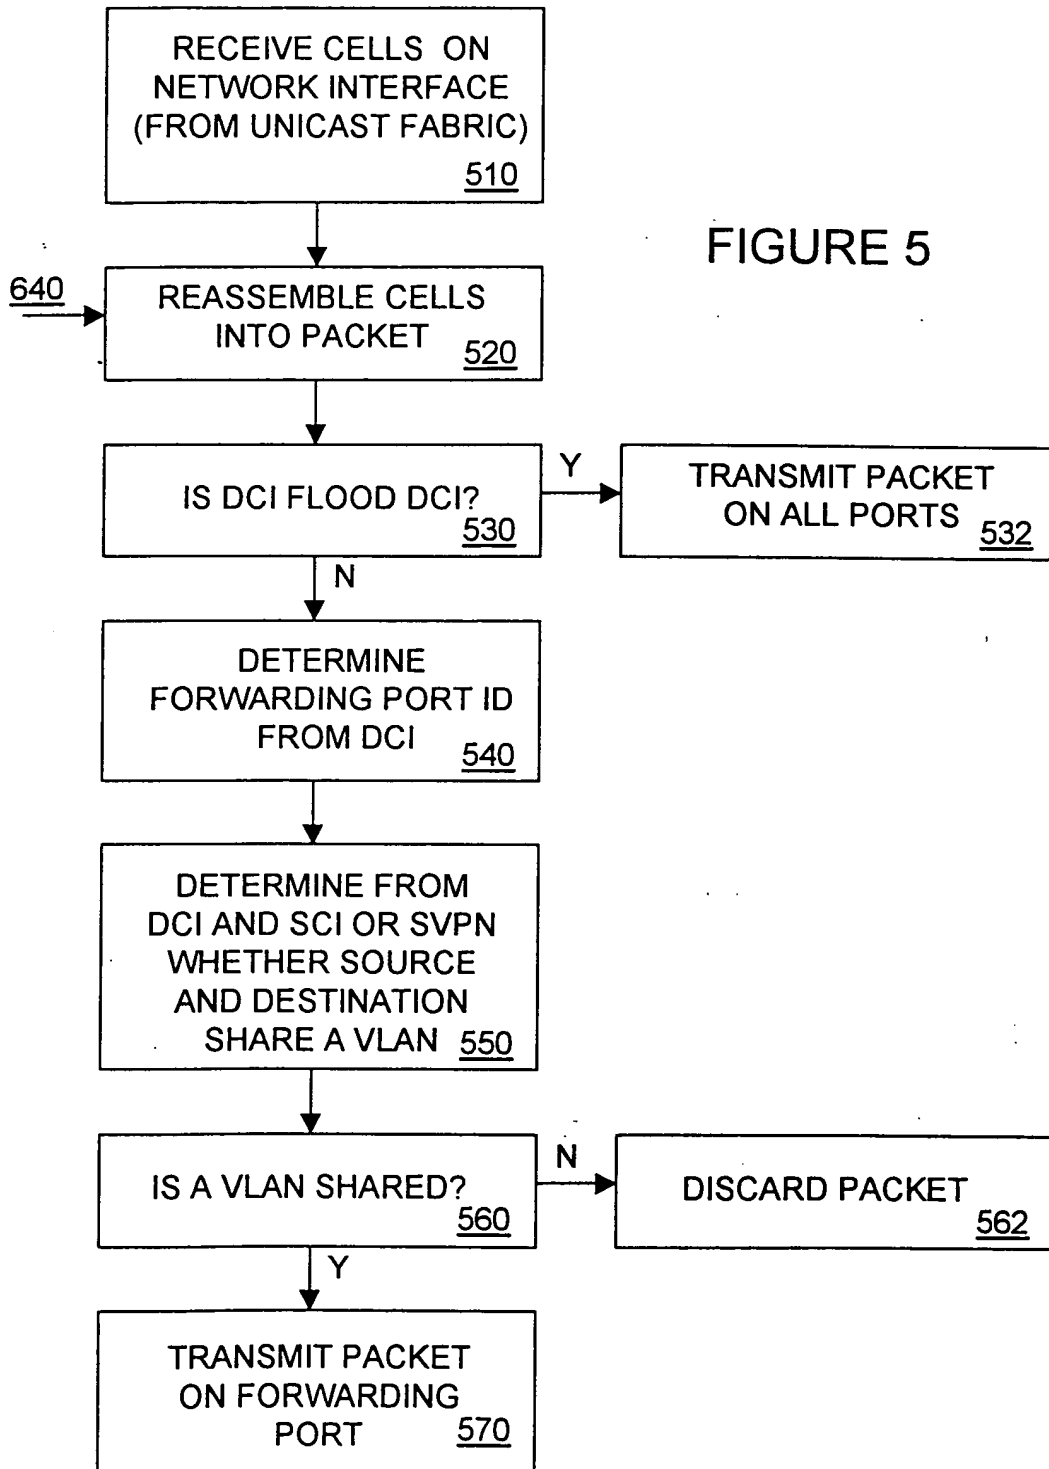


ALCATEL CONFIDENTIAL

FIGURE 4



ALCATEL CONFIDENTIAL



ALCATEL CONFIDENTIAL

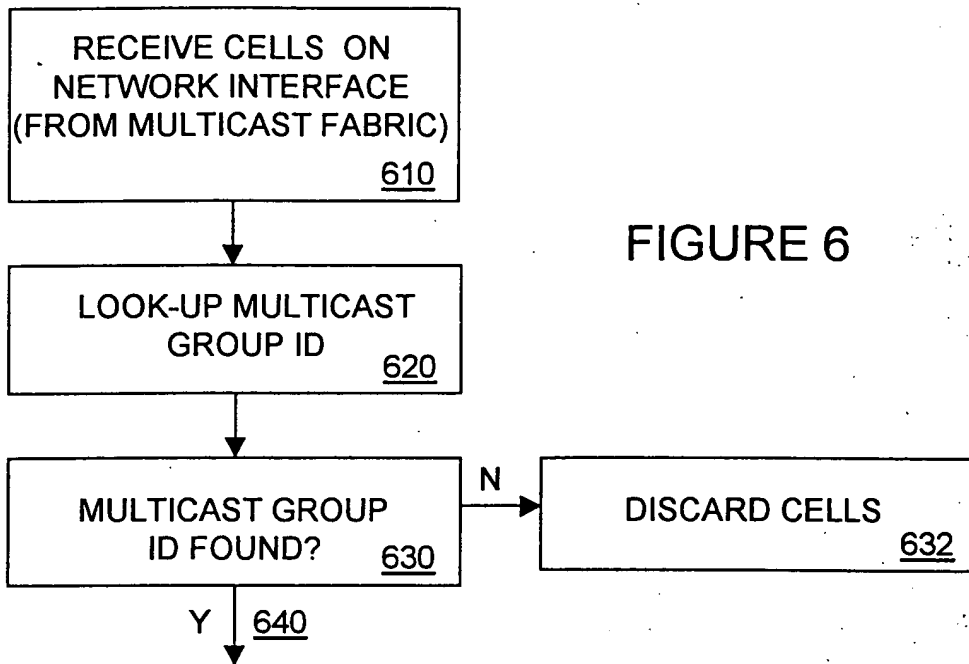
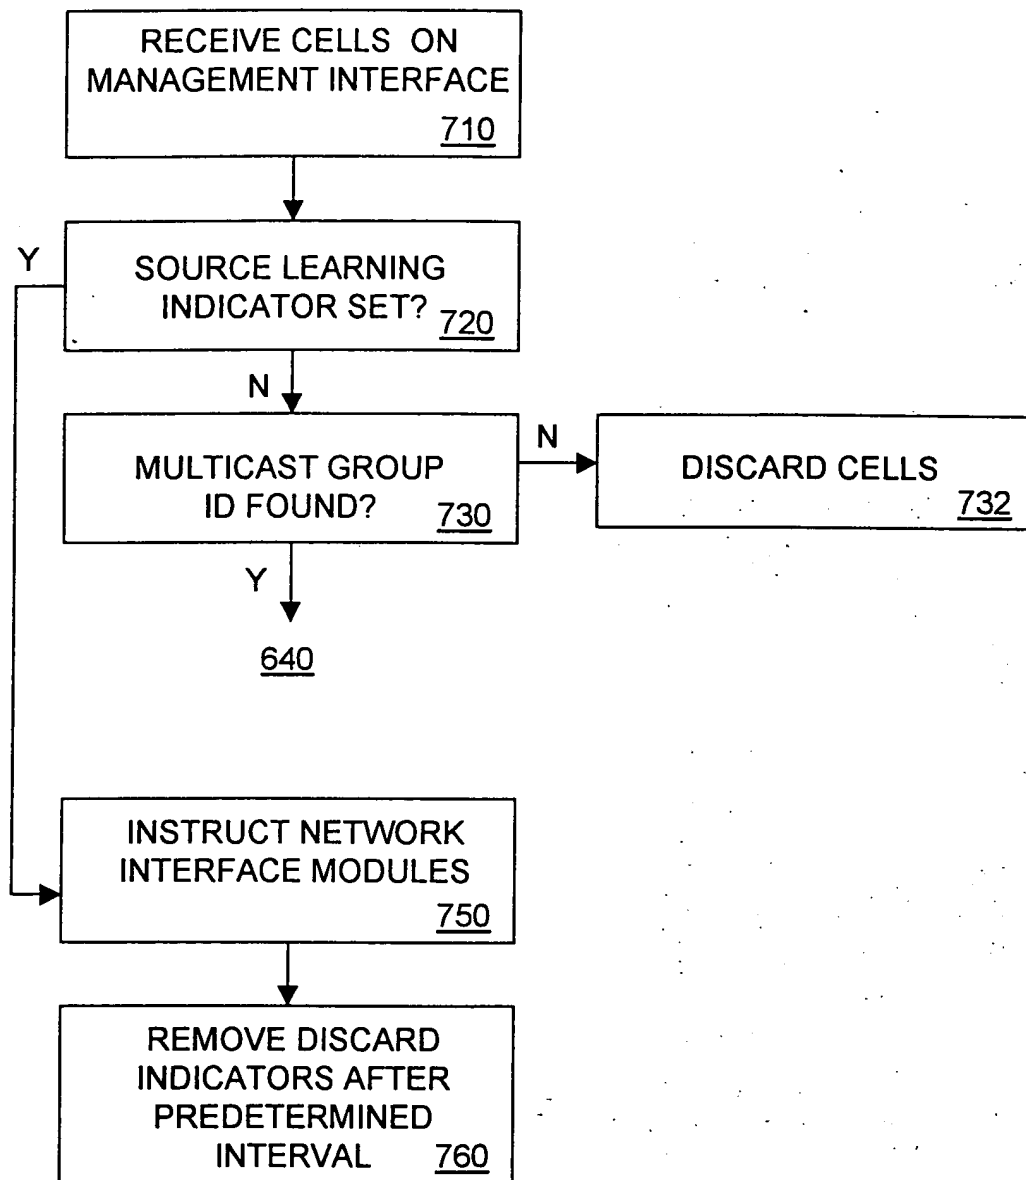


FIGURE 6

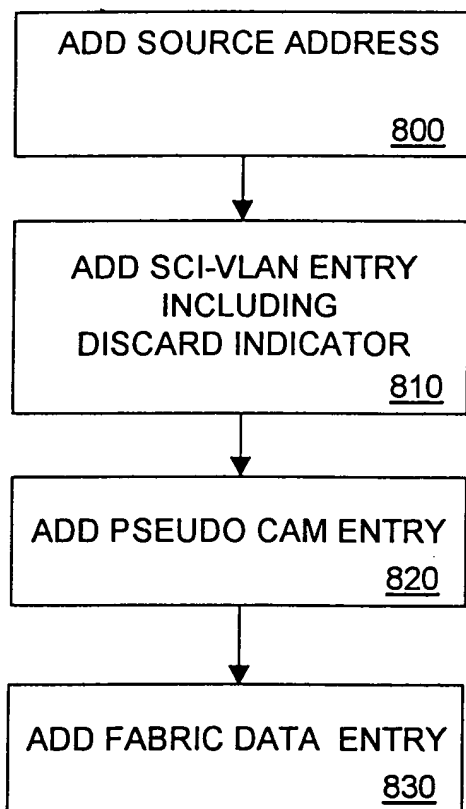
ALCATEL CONFIDENTIAL




ALCATEL CONFIDENTIAL

FIGURE 7

**INSTRUCTIONS TO THE NETWORK
INTERFACE MODULE THAT
INTERFACES WITH THE SOURCE**



**INSTRUCTIONS TO OTHER
NETWORK INTERFACE MODULES**



ALCATEL CONFIDENTIAL

FIGURE 8

EXHIBIT D

FedEx Tracking Number

8182 6492 6623

From *Please print and press hard.* 1996-4434-3
 Date 4/19/01 Sender's FedEx Account Number
 Sender's Name D. Cordeiro Phone (818) 878 5080

Company ALCATEL INTERNETWORKING

Address 26801 W AGOURA RD

City CALABASAS State CA ZIP 91301

Your Internal Billing Reference OPTIONAL

To Recipient's Name Eugene L. Hahn Phone (650) 853 3600
650 470 2900

Company Coudert Brothers

Address 530 Lytton Avenue

We cannot deliver to P.O. boxes or P.O. ZIP codes.

Suite 300

To "HOLD" at FedEx location, print FedEx address here.

City Palo Alto State CA ZIP 94301

NEW Peel and Stick FedEx USA Airbill

See back for application instructions.

Questions? Call 1-800-Go-FedEx® (800-463-3339)

Visit our Web site at www.fedex.com

By using this Airbill you agree to the service conditions on the back of this Airbill and in our current Service Guide, including terms that limit our liability.

0124882269

4a Express Package Service *Packages up to 150 lbs. Delivery commitment may be later in some areas.*
☐ FedEx Priority Overnight Next business morning
☒ FedEx Standard Overnight Next business afternoon
☐ FedEx First Overnight Earliest next business morning delivery to select locations
☐ FedEx 2Day* Second business day
☐ FedEx Express Saver* Third business day
 * FedEx Letter Rate not available Minimum charge: One-pound rate

4b Express Freight Service *Packages over 150 lbs. Delivery commitment may be later in some areas.*
☐ FedEx 1Day Freight* Next business day
☐ FedEx 2Day Freight Second business day
☐ FedEx 3Day Freight Third business day
 * Call for Confirmation

5 Packaging ** Declared value limit \$500*
☒ FedEx Letter* ☐ FedEx Pak* ☐ Other Pkg. Includes FedEx Box, FedEx Tube, and customer pkg.

6 Special Handling
☐ Saturday Delivery Available for FedEx Priority Overnight and FedEx 2Day to select ZIP codes
☐ Sunday Delivery Available for FedEx Priority Overnight to select ZIP codes
☐ HOLD Weekday at FedEx Location Not available with FedEx First Overnight
☐ HOLD Saturday at FedEx Location Available for FedEx Priority Overnight and FedEx 2Day to select locations
 Does this shipment contain dangerous goods? *Use box must be checked.*
☒ No ☐ Yes As per attached Shipper's Declaration ☐ Yes Shipper's Declaration not required
☐ Dry Ice Dry Ice, 9 UN 1845 x _____ kg
 Dangerous Goods cannot be shipped in FedEx packaging. ☐ Cargo Aircraft Only

7 Payment *Bill to:* Enter FedEx Acct. No. or the bill Card No. below.
☐ Sender Acct. No. in Section 1 will be billed. ☐ Recipient ☐ Third Party ☐ Credit Card ☐ Cash/Check

FedEx Acct. No. Credit Card No. Exp. Date
 Total Packages Total Weight Total Declared Value*
 \$.00
 * Our liability is limited to \$100 unless you declare a higher value. See back for details. FedEx Use Only

8 Release Signature *Sign to authorize delivery without obtaining signature.*

By signing you authorize us to deliver this shipment without obtaining a signature and agree to indemnify and hold us harmless from any resulting claims.

359

SPR 1199 - Rev. Date 11/98 - Part #1548135 - © 1994-98 FedEx - PRINTED IN U.S.A.

FedEx Tracking Number

8182 6492 6597

From *Please print and press hard.* 1996-4434-3
 Date 2/16/01 Sender's FedEx Account Number
 Sender's Name Scot Reader Phone (818) 878 4799

Company ALCATEL INTERNETWORKING

Address 26801 W AGOURA RD

City CALABASAS State CA ZIP 91301

Your Internal Billing Reference OPTIONAL

To Recipient's Name Eugene L. Hahn Phone (650) 853 3600
650 470 2900

Company Coudert Brothers

Address 530 Lytton Avenue

We cannot deliver to P.O. boxes or P.O. ZIP codes.

Suite 300

To "HOLD" at FedEx location, print FedEx address here.

City Palo Alto State CA ZIP 94301-1541

NEW Peel and Stick FedEx USA Airbill

See back for application instructions.

Questions? Call 1-800-Go-FedEx® (800-463-3339)

Visit our Web site at www.fedex.com

By using this Airbill you agree to the service conditions on the back of this Airbill and in our current Service Guide, including terms that limit our liability.

0124882269

4a Express Package Service *Packages up to 150 lbs. Delivery commitment may be later in some areas.*
☐ FedEx Priority Overnight Next business morning
☒ FedEx Standard Overnight Next business afternoon
☐ FedEx First Overnight Earliest next business morning delivery to select locations
☐ FedEx 2Day* Second business day
☐ FedEx Express Saver* Third business day
 * FedEx Letter Rate not available Minimum charge: One-pound rate

4b Express Freight Service *Packages over 150 lbs. Delivery commitment may be later in some areas.*
☐ FedEx 1Day Freight* Next business day
☐ FedEx 2Day Freight Second business day
☐ FedEx 3Day Freight Third business day
 * Call for Confirmation

5 Packaging ** Declared value limit \$500*
☐ FedEx Letter* ☐ FedEx Pak* ☒ Other Pkg. Includes FedEx Box, FedEx Tube, and customer pkg.

6 Special Handling
☐ Saturday Delivery Available for FedEx Priority Overnight and FedEx 2Day to select ZIP codes
☐ Sunday Delivery Available for FedEx Priority Overnight to select ZIP codes
☐ HOLD Weekday at FedEx Location Not available with FedEx First Overnight
☐ HOLD Saturday at FedEx Location Available for FedEx Priority Overnight and FedEx 2Day to select locations
 Does this shipment contain dangerous goods? *Use box must be checked.*
☒ No ☐ Yes As per attached Shipper's Declaration ☐ Yes Shipper's Declaration not required
☐ Dry Ice Dry Ice, 9 UN 1845 x _____ kg
 Dangerous Goods cannot be shipped in FedEx packaging. ☐ Cargo Aircraft Only

7 Payment *Bill to:* Enter FedEx Acct. No. or the bill Card No. below.
☒ Sender Acct. No. in Section 1 will be billed. ☐ Recipient ☐ Third Party ☐ Credit Card ☐ Cash/Check

FedEx Acct. No. Credit Card No. Exp. Date
 Total Packages Total Weight Total Declared Value*
 \$.00
 * Our liability is limited to \$100 unless you declare a higher value. See back for details. FedEx Use Only

8 Release Signature *Sign to authorize delivery without obtaining signature.*

By signing you authorize us to deliver this shipment without obtaining a signature and agree to indemnify and hold us harmless from any resulting claims.

359

EXHIBIT E

Subject: IMC

Date: Mon, 16 Apr 2001 14:41:00 -0700

From: Dave Cordeiro <Dave.Cordeiro@ind.alcatel.com>

Organization: Alcatel-eND

To: ehahm@coudert.com

Dear Mr. Hahm,

Myself and Mr. Scot Reader sent you some documents on February 16, 2001, to be signed by your clients Christopher Hoogenboom and John Wallner of IMC. The enclosed documents were:

I. U.S. Patent Application No. 09/718,046 entitled "Source Learning for Multi-Path Backplane Data Switch" and its respective Declaration and Power of Attorney and Assignment forms;

II. U.S. Patent Application No. 09/718,050, entitled "Flow Integrity for Path Transitioning Data Switch" and its respective Declaration and Power of Attorney and Assignment forms; and

III. U.S. Patent Application No. 09/098,228, entitled "Digital Traffic Switch with Credit-Based Buffer Control" and its respective Assignment form.

Could you please tell me if your clients intend to execute these documents or not?

Sincerely,
David A. Cordeiro
IP Counsel

--
This email may contain confidential and privileged material for the sole use of the intended recipient(s). Any review, use, distribution or disclosure by others is strictly prohibited. If you are not the intended recipient (or authorized to receive for the recipient), please contact the sender by reply email and delete all copies of this message.

David Cordeiro <dave.cordeiro@ind.alcatel.com>

Intellectual Property Counsel
e-Business Networking Division
Intellectual Property

EXHIBIT F

Mr. Eugene L. Hahm
Coudert Brothers
530 Lytton Avenue
Suite 300
Palo Alto, CA 94301

Dear Mr. Hahm:

We have attempted to contact you several times regarding your clients, Chris Hoogenboom and John Wallner. We believe that Chris Hoogenboom and John Wallner are joint applicants of inventions 134001 and 134002, entitled FLOW INTEGRITY FOR PATH TRANSITIONING DATA SWITCH and SOURCE LEARNING FOR MULTI-PATH BACKPLANE DATA SWITCH, respectively. In addition, we believe Chris Hoogenboom to be a joint applicant of invention 111542, entitled DIGITAL TRAFFIC SWITCH WITH CREDIT-BASED BUFFER CONTROL.

If Chris Hoogenboom and John Wallner are inventors, as discussed above, they are obligated to by their past employment agreements with Alcatel to sign the Declaration and Assignment and Power of Attorney forms which were sent to you on February 16, 2001. Please contact your clients, Chris Hoogenboom and John Wallner, so that they can fulfill their duty to sign the Declaration and Assignment and Power of Attorney forms for the subject above mentioned inventions of which you have in your possession. If you have any questions, please feel free to contact me.

Sincerely,

David A. Cordeiro
Patent Counsel

EXHIBIT G

FedEx USA Airbill

FedEx Tracking Number

8182 6492 6623

From Please print and press hard. Sender's FedEx Account Number 1996-4434-3
 Date 4/19/01
 Sender's Name D. Cordero Phone (818) 878 5080

Company ALCATEL INTERNETWORKING

Address 26801 W AGOURA RD

City CALABASAS State CA ZIP 91301

Your Internal Billing Reference
 First 24 characters will appear on invoice.

To Recipient's Name Eugene L. Hahn Phone (650) 853 3600
 650 470 2900

Company Coudert Brothers

Address 530 Lytton Avenue

We cannot deliver to P.O. boxes or P.O. ZIP codes.

To "HOLD" at FedEx location, print FedEx address here.

City Palo Alto State CA ZIP 94301

NEW Paper and Stock FedEx USA Airbill

See back for application instructions.

Questions? Call 1-800-Go-FedEx® (800-463-3339)

Visit our Web site at www.fedex.com

By using this Airbill you agree to the service conditions on the back of this Airbill and in our current Service Guide, including terms that limit our liability.

0124882269

SLA 1515
 0215
 Sender's Copy

4a Express Package Service
☐ FedEx Priority Overnight Next business morning
☒ FedEx Standard Overnight Next business afternoon
☐ FedEx First Overnight Earliest next business morning delivery to select locations
 Packages up to 150 lbs. Delivery commitment may be later in some areas.

☐ FedEx 2Day® Second business day
☐ FedEx Express Saver® Third business day
 * FedEx Letter Rate not available Minimum charge: One-pound rate

4b Express Freight Service
☐ FedEx 10day Freight® Next business day
☐ FedEx 2day Freight Second business day
☐ FedEx 3day Freight Third business day
 Packages over 150 lbs. Delivery commitment may be later in some areas.

* Call for Confirmation

5 Packaging
☒ FedEx Letter®
☐ FedEx Pak®
☐ Other Pkg. Includes FedEx Box, FedEx Tube, and customer pkg.
 * Declared value limit \$500

6 Special Handling
☐ Saturday Delivery Available for FedEx Priority Overnight and FedEx 2Day to select ZIP codes
☐ Sunday Delivery Available for FedEx Priority Overnight to select ZIP codes
☐ HOLD Weekday at FedEx Location Not available with FedEx First Overnight
☐ HOLD Saturday at FedEx Location Available for FedEx Priority Overnight and FedEx 2Day to select locations

Does this shipment contain dangerous goods?
 Line 104 must be checked.
☒ No ☐ Yes As per attached Shipper's Declaration
☐ Yes Shipper's Declaration not required
☐ Dry Ice Dry Ice, 3 UN 1845 x kg
☐ Cargo Aircraft Only
 Dangerous Goods cannot be shipped in FedEx packaging.

7 Payment Bill to:
☐ Sender Acct. No. in Section 1 will be billed.
☐ Recipient
☐ Third Party
☐ Credit Card
☐ Cash/Check
 Enter FedEx Acct. No. or bill to Card No. below.

FedEx Acct. No. Credit Card No. Exp. Date
 Total Packages Total Weight Total Declared Value*
 \$.00
 *Our liability is limited to \$100 unless you declare a higher value. See back for details. FedEx Use Only

8 Release Signature Sign to authorize delivery without obtaining signature.

By signing you authorize us to deliver this shipment without obtaining a signature and agree to indemnify and hold us harmless from any resulting claims.

359

SRP 1199-Rev. 04/98-Part #1548135-© 1994-98 FedEx-PRINTED IN U.S.A.

FedEx USA Airbill

FedEx Tracking Number

8182 6492 6597

From Please print and press hard. Sender's FedEx Account Number 1996-4434-3
 Date 2/16/01
 Sender's Name Scot Reader Phone (818) 878 4799

Company ALCATEL INTERNETWORKING

Address 26801 W AGOURA RD

City CALABASAS State CA ZIP 91301

Your Internal Billing Reference
 First 24 characters will appear on invoice.

To Recipient's Name Eugene L. Hahn Phone (650) 853 3600
 650 470 2900

Company Coudert Brothers

Address 530 Lytton Avenue

We cannot deliver to P.O. boxes or P.O. ZIP codes.

To "HOLD" at FedEx location, print FedEx address here.

City Palo Alto State CA ZIP 94301-1541

NEW Paper and Stock FedEx USA Airbill

See back for application instructions.

Questions? Call 1-800-Go-FedEx® (800-463-3339)

Visit our Web site at www.fedex.com

By using this Airbill you agree to the service conditions on the back of this Airbill and in our current Service Guide, including terms that limit our liability.

SLA 1515
 0215
 Sender's Copy

4a Express Package Service
☐ FedEx Priority Overnight Next business morning
☒ FedEx Standard Overnight Next business afternoon
☐ FedEx First Overnight Earliest next business morning delivery to select locations
 Packages up to 150 lbs. Delivery commitment may be later in some areas.

☐ FedEx 2Day® Second business day
☐ FedEx Express Saver® Third business day
 * FedEx Letter Rate not available Minimum charge: One-pound rate

4b Express Freight Service
☐ FedEx 10day Freight® Next business day
☐ FedEx 2day Freight Second business day
☐ FedEx 3day Freight Third business day
 Packages over 150 lbs. Delivery commitment may be later in some areas.

* Call for Confirmation

5 Packaging
☐ FedEx Letter®
☐ FedEx Pak®
☒ Other Pkg. Includes FedEx Box, FedEx Tube, and customer pkg.
 * Declared value limit \$500

6 Special Handling
☐ Saturday Delivery Available for FedEx Priority Overnight and FedEx 2Day to select ZIP codes
☐ Sunday Delivery Available for FedEx Priority Overnight to select ZIP codes
☐ HOLD Weekday at FedEx Location Not available with FedEx First Overnight
☐ HOLD Saturday at FedEx Location Available for FedEx Priority Overnight and FedEx 2Day to select locations

Does this shipment contain dangerous goods?
 Line 104 must be checked.
☒ No ☐ Yes As per attached Shipper's Declaration
☐ Yes Shipper's Declaration not required
☐ Dry Ice Dry Ice, 3 UN 1845 x kg
☐ Cargo Aircraft Only
 Dangerous Goods cannot be shipped in FedEx packaging.

7 Payment Bill to:
☒ Sender Acct. No. in Section 1 will be billed.
☐ Recipient
☐ Third Party
☐ Credit Card
☐ Cash/Check
 Enter FedEx Acct. No. or bill to Card No. below.

FedEx Acct. No. Credit Card No. Exp. Date
 Total Packages Total Weight Total Declared Value*
 \$.00
 *Our liability is limited to \$100 unless you declare a higher value. See back for details. FedEx Use Only

8 Release Signature Sign to authorize delivery without obtaining signature.

By signing you authorize us to deliver this shipment without obtaining a signature and agree to indemnify and hold us harmless from any resulting claims.

35